Does Electronic Filing and Payment Increase Tax Compliance? Evidence from Large Taxpayers in Senegal *

Léo Czajka Awa Diouf Fabrizio Santoro

7th July, 2026

Abstract

Governments in low-income countries have progressively introduced electronic tax filing and payment systems in the hope of reducing enforcement costs for tax administrations, compliance costs for taxpayers, and the risk of collusive in-person interactions between the two. Combining high-frequency administrative data with a dynamic difference-in-differences approach, we investigate the causal impact of a reform that made the use of these technologies mandatory for large taxpayers in Senegal. Our findings indicate no—or only limited—effects on key measures of tax compliance, such as the probability of declaring, the probability of paying, or tax payments. However, e-filing reduces the prevalence of missing values by more than 90 per cent, with notable measurement implications. In particular, we show that aggregate formal employment is at least 20 per cent greater than suggested by digitised paper-based records.

Keywords— O33 – Technological Change: Choices and Consequences; Diffusion Processes, D73 – Bureaucracy; Administrative Processes in Public Organizations; H26 – Tax Evasion and Avoidance

*Léo Czajka: International Tax Observatory – Paris School of Economics; Awa Diouf and Fabrizio Santoro: International Centre for Tax and Development (ICTD). The authors gratefully acknowledge financial support from the Norwegian Agency for Development Cooperation (NORAD grant QZA-22/0011), the European Union (TAXUD/2023/DE/2018) and the International centre for Tax and Development. We thank all members of the *Direction Générale des Impôts et Domaines* (DGID) who we have been working with. In particular we thank the general director Abdoulaye Diagne, the former general directors Bassirou Samba Niasse and Cheikh Ahmed Tidiane Bâ, the technical advisors Mor Fall and Sérigne Mabeye Fall, the director of the intelligence department Jean Koné, the director of the large firms directorate Hady Dieye, the director of the medium firms directorate Aïssatou Ndao, the former director of the IT department Yama Kouyate Diaby. We also thank Aliou Sow and Mamadou Guisse for their administrative support over the years, and IT experts, Oumar Ngom and Mamadou Diangar for their technical support. All errors are our own.

1. Introduction

Efficient institutions are widely regarded as central pillars of economic development. In this regard, the accelerating production and global diffusion of information technologies has the potential to enhance the performance of public administrations in low-income countries (LICs), and thus contribute to their convergence toward higher living standards. In recent decades, LICs have heavily invested in the adoption of information technology infrastructure[1] and services, including online platforms for declaring taxes electronically (e-filing) and paying taxes electronically (e-payment).[2] In contexts where domestic revenues remain insufficient to adequately finance essential public goods (OCDE, 2022), such tools are expected to contribute to higher tax revenues.

The expectations regarding e-filing and e-payment rest on several mechanisms. First, these technologies may reduce enforcement costs for tax administrations by reducing paperwork and the cost of accessing information from taxpayers. Second, they should reduce compliance costs for the taxpayers at least by saving them transportation costs. Third, by minimising in-person interactions between taxpayers and tax officials, they may limit opportunities for collusion or extortion. However, these impacts may not translate into improved compliance. On the taxpayers' side, cost reductions may be allocated to purposes other than better complying with taxes. Administrations may decide to allocate newly available resources to functions that are not related to enforcement. The impact of reduced in-person interactions depends on the prevalence of collusive practices ex-ante (Okunogbe and Pouliquen, 2022) and on how binding the technology is. Therefore, the effect of electronic taxation on tax compliance remains an empirical question.

This article provides causal evidence on the impact of the joint adoption of e-filing and e-payment in Senegal. This technology (Etax) was made mandatory by ministerial decree in

[1] See GovTech Maturity Index (GTMI) Data Dashboard, (accessed 6 March 2026).
[2] As of 2016, 32 per cent of LICs had adopted e-filing (World Bank, 2016)

July 2017 for all taxpayers monitored by the Large Taxpayers Centre (LTC). By January 2018, approximately 80 per cent of targeted taxpayers had registered on and were using the platform. Using newly collected administrative data from three monthly taxes—value-added tax (VAT), withholding tax on wages (or "Pay As You Earn" taxes (PAYE)), and withholding tax on service suppliers (WTS)—combined with a dynamic difference-in-differences design, we exploit exogenous variation in technology adoption generated by the reform to estimate its causal impact on multiple indicators of tax compliance. To improve the comparability across groups, we restrict the control group to the largest taxpayers outside the LTU. While our setting presents some identification challenges, it has great external validity, as it closely mirrors the way tax administrations typically deploy new IT systems: by starting with the largest, most closely monitored taxpayers and gradually expanding to smaller firms (Ali et al., 2015). Our main specification uses a strictly balanced panel of 1,400 taxpayers observed monthly over the period 2017–2019. We impose balance to avoid confounding real effects with compositional changes.

Our main result is a sharp and highly persistent decline of approximately 6-7 percentage points in the probability that a tax declaration has a missing value. Relative to pre-reform average, this represents a 90 per cent decrease. We interpret this effect as a change in the data-production process rather than in taxpayer behaviour. Prior to e-filing, declarations were submitted on paper and digitised by data-entry agents, who had limited incentives to complete all fields not strictly required to compute tax liability. By allowing taxpayers to enter information directly into the system, e-filing eliminates these omissions and mechanically reduces the prevalence of missing values. We show that this improvement has important measurement implications. In particular, the field 'number of employees' in monthly PAYE declarations becomes substantially more complete. As a result, the average number of employees reported by treated firms increases immediately by more than 40 workers per firm on average. This represents a 33 per cent rise relative to the pre-reform level. The magnitude of this increase alone accounts for roughly two-thirds of the sudden 30 per cent rise in aggregate formal employment observed at the national level between September and December 2017. Altogether, this implies that, because of missing values, aggregate formal employment was underestimated by at least 20 per cent before the adoption of e-filing.

We find limited or insignificant results on other margins of compliance. We document a short-lived decline of 2–3 percentage points in the probability of declaring immediately after adoption. This is probably a result of transition costs faced both by taxpayers adjusting to a new technology and by the tax administration rapidly registering a large number of firms and adapting established manual procedures. We also find no clear evidence that the reform increased tax collection. Tax payments did not increase at the extensive margin. At the intensive margin, we do observe an upward trend, but *i)* the magnitude of this increase is not sufficient to rule out underlying time trends *ii)* we show that at least 30 per cent of the observed increase in tax payments among treated firms can be accounted for by improvements in the digital recording of payments, rather than by genuine increases in tax collection.

This paper contributes to the still scarce literature on the role of e-filing and e-payment technology on taxation in LICs (Okunogbe and Santoro, 2023). To our knowledge, there exist only three studies exploring the causal impact of e-filing and e-payment on tax compliance with a counterfactual method in a low/middle income country (Jouste, Nalukwago, and Waiswa, 2021 in Uganda, Okunogbe and Pouliquen, 2022 for Tajikistan, Santoro, Amine, and Magongo, 2024 for Eswatini). Our study differs from the latter primarily in that it investigates the behaviour of the largest taxpayers. These taxpayers are by far the most strategic sub-sample of taxpayers when considering total tax revenue (Milanez, 2017, Slemrod and Velayudhan, 2018). Our study therefore speaks more directly to the aggregate implications of adopting such technologies than do previous studies. Second, our use of high-frequency monthly data allows us to precisely and more credibly document the evolution of treatment impact in the short and mid-term.

Overall, our evidence suggests that the adoption of e-taxation has no—or only limited—effects on tax compliance in the short to medium term. Our findings illustrate that the expected returns to technology adoption by government institutions may not materialise immediately, even in contexts where adoption is mandatory. One plausible explanation for the absence of detectable effects is that the policy was implemented exclusively among large taxpayers, who tend to be more compliant ex ante than smaller firms (Almunia and Lopez-Rodriguez, 2018, Bachas, Jaef, and Jensen, 2019). But, since large taxpayers typically account for most tax revenue, our findings imply that the adoption of these technologies is unlikely to produce large aggregate revenue gains, at least in the short- to medium-term. On the other hand, we show that e-taxation can substantially improve the quality of microdata produced by tax administrations. In particular, our results indicate that, in the absence of online taxation, key statistics such as aggregate formal employment can be significantly underestimated. This finding is especially relevant for low-income countries, where weak statistical capacity systematically undermines governments' ability to identify problems and design effective public policies. From a methodological standpoint, our results also offer a cautionary note for the growing body of academic research that relies on administrative data from countries that are undergoing digital transformation.

The rest of the paper is structured as follows: section 2 documents the context; section 3 specifies our assumptions on taxpayers and administrators' behaviour as well as the main expected channels; section 4 describes our data and empirical strategy; section 5 discusses the results; and section 6 concludes.

2. Institutional Context

The Senegalese Tax Administration (DGID) directly monitors about 30,000 taxpayers (mostly firms) distributed across 20 tax centres according to business size, sector, and location. Taxpayers with a turnover exceeding USD 3,200,000 for three consecutive years are managed by the large taxpayers centre (LTC), irrespective of their headquarters' location. Smaller taxpayers operating in Dakar are distributed across two medium-sized taxpayers centres (MTC1, MTC2), one unit for firms operating in regulated sectors (RSC) and seven municipal tax centres.[3] The rest of the taxpayers are spread across regional centres.

In late 2016, the DGID started implementing and testing new digital software (Etax) allowing taxpayers to declare online (e-filing) and pay online (e-payment). The 36 largest taxpayers were progressively onboarded during this pilot phase, and in July 2017 the Senegalese Minister of Economy and Finance released a decree making it mandatory for all taxpayers monitored by the LTC to use Etax. The adoption of Etax was subsequently encouraged among taxpayers from other centres (MTC1, MTC2 and RSC), but with limited take-up. Before the adoption of Etax at the LTC, tax declarations were submitted in paper format, mostly by taxpayers themselves and not by mail, and then digitised into the system by data-entry agents. Taxpayers would also have to come to their tax centre in order to pay (mostly in cash or checks). The reform thus had multiple official objectives: *i)* to spare taxpayers the cost of moving to their centre to declare and/or pay, *ii)* to spare the administration the cost of handling all paper declarations, and *iii)* to improve the quality of the information by removing mistakes introduced by data-entry agents.

E-filing through Etax mimics a paper declaration as closely as possible: only identification information is pre-filled, and inconsistencies as well as missing values are allowed and not automatically flagged. There is no automatic reminder for missing declarations. With a few minor exceptions, all types of taxes can be declared through Etax. Tax liabilities are automatically computed by the system based on the information reported by the taxpayer. If a taxpayer realises they have made a mistake, they must submit a correction request which is then handled by the tax centre. Taxpayers can consult their data at all times. E-payment occurs through the banking system.

[3]Regulated sectors are medical and legal professions, accountants, architects, state agencies and NGOs. Firms within the MTC1 (MTC2) have a turnover between USD 1,600,000 and USD 3,200,000 (USD 800,000 and USD 1,600,000). Sector prevails over size for fiscal assignment. Dakar's seven municipal centres are Dakar-Plateau, Ngor-Almadies, Parcelles Assainies, Dakar Liberté, Grand Dakar, Pikine-Guédiawaye and Rufisque.

3. Assumptions and Expected Impacts

General assumptions In accordance with the classic deterrence model for tax evasion first developed by Allingham and Sandmo, 1972 as well as a large corpus of empirical literature (see Luttmer and Singhal, 2014 and Slemrod, 2019), we assume that taxpayers mostly act as free-riders. In other words, in the absence of any risk associated with tax evasion, economic actors would simply not pay taxes. In addition, following a rich corpus of work that builds from the seminal work of Weber, 1922, we assume that the principal-agent model is an adequate framework to reflect on the behaviour of bureaucrats. Here, the principal will typically be a higher-tier bureaucrat, and the agents will be civil servants under their supervision, in charge of performing tasks contributing to tax collection, the final objective of the tax administration.[4]

Expected impacts: We list here what we identify as the three main channels through which e-filing and e-payment can affect tax compliance. We formulate hypotheses on their effects, based on our understanding of the context, the technology's specifics, and the general assumptions set out above. The objective of this section is to explain more precisely why and how the mandatory adoption of e-taxation by the largest taxpayers in Senegal could impact their compliance, but we must stress that the available data does not allow us to disentangle each mechanism.

First channel - reduction of compliance costs: The possibility of declaring and paying online allows taxpayers to save travel costs. According to the free-rider taxpayers model, this should increase the probability of declaring or paying only for taxpayers who, ex-ante, were not compliant because of travel costs. However, in the specific setting of this study, where the treated taxpayers are the largest firms of the country, it is reasonable to assume that travel costs would not be such a binding constraint. We therefore assume that this channel should play little role in explaining ex-post variation in compliance.

Second channel - reduction of enforcement costs: Before the introduction of Etax, all declarations

[4]See Besley et al., 2022 for a formal exposition of this general model.

were submitted in paper format (see section 2). The dramatic reduction in the cost of processing tax declarations has two main implications.

1. Information becomes more reliable. E-filing indeed eliminates mistakes introduced during data-entry of paper format declarations. In addition, under e-filing, a missing declaration straightforwardly indicates that the taxpayers failed to declare, while, before e-filing, it could have been lost (or not yet digitised). Better information facilitates tax collectors' job, which could translate into improved compliance.

2. E-filing inevitably frees up some time for agents who previously worked on paper processing. Such a gain will inevitably be observed by these agents' principals, and we should expect the agents to be given new tasks. However, the impact on tax compliance is ambiguous, as it will depend on the nature of such new tasks.

Third channel - limiting in-person interactions: In-person interactions between bureaucrats and taxpayers may affect tax collection in opposite ways. In one way, they can collude to the detriment of the government: taxpayers may reduce their effective liability in exchange for paying the bureaucrat a share of the implied gain. In that case, in-person interactions reduce tax revenue. Inversely, tax officials may also force taxpayers to pay more than they owe to the state in order to meet their own tax revenue target. Suggestive evidence for both channels has been found by Okunogbe and Pouliquen, 2022. However, in the context of this study, given that treated taxpayers are very large, we expect the bargaining power to be in their favour when interacting in person with a tax official. Hence the ex-ante prevalence of overpayments should not be empirically relevant here, and we expect the adoption of Etax to increase tax payment conditional on ex-ante collusion.

4. Data and Empirical Strategy

4.1. Data

We compile four datasets, all generated by the tax administration as part of its routine operations. In addition to the source specific information listed below, each declaration/payment comes with an identification number, the name of the centre it was recorded by, and whether it was processed through Etax or not. We use this information to identify taxpayers' fiscal centre and trace their effective use of the e-filing and e-payment.

- *Value-added tax (VAT)*: monthly, taxpayer level, includes information on total sales, VAT due and information used to compute the VAT due.
- *Withholding taxes on employees (PAYE)*: monthly, taxpayer level, includes information about the number of employees, total wage bill and total liability declared.
- *Withholding taxes on informal service suppliers (WTS)*: monthly, taxpayers level, includes information on the total amount spent in exchange for services from informal suppliers, the total liability declared.
- *Payment data*: daily, transaction level, includes information on the amount paid, the nature of the tax, the fiscal period it applies to, and the mode of payment (cash, check, electronic, etc.).

Tax compliance encompasses multiple dimensions. For the taxpayers, it implies declaring true and complete information on time, and paying all the taxes due on time. On the administrative side, it implies accurately recording the information submitted by the taxpayers, computing the correct liability and recovering the right amount of taxes. Considering this, the probability that a declaration is recorded in the system captures the form of compliance that is easiest to measure. Taxpayers have to declare even when their liability is null, so the absence of declaration implies either that the taxpayers did not declare, or—in the case of a paper format declaration—that the administration did not digitise the declaration. Other forms of compliance are harder to measure because they rely on information that can be subject to errors and manipulations. Indeed, using only administrative data on tax declarations and payments, we cannot know with certainty how much taxpayers ought to declare, nor whether and how much they ought to pay. Nevertheless, if we assume that taxpayers prefer not to pay taxes (see section 3), any observed increase in indicators implying that taxpayers (could) pay more taxes, is indicative of improvement in tax compliance unless it stems from an increase in real activity or extortion by tax officials (unlikely in our setting, see section 3).

We thus estimate our model on five main outcomes: a) the probability of declaring; b) the liability declared; c) the probability of paying; d) the amount paid; e) the probability of a value being missing. In the appendix, we also report graphical results on f) total liability declared, outstanding liability and outstanding liability as a percentage of total liability. Each variable is measured by aggregating information from the three monthly taxes listed above. We consider a declaration has a *missing* value if a positive liability is associated with the declaration but at least one variable that should be positive (e.g. total wage bill for PAYE) is missing or equal to 0.

4.2. Sample selection

For each of the three tax specific databases, we restrict to taxpayers who have declared and/or paid at least once before January 2017 *and* beyond December 2019. We only keep taxpayers that we know existed and were subject to the tax before and after this time window. Due to missing declarations, this inevitably creates unbalanced panels. We 'fill the gaps' by adding observations for every missing declaration between January 2017 and December 2019. For these observations, all values are set to zero. Importantly, this implies that missing declarations have no missing values: our definition of missing value is, by construction, conditional on having declared. Next, we aggregate the three balanced panels at the taxpayer-month level.

We cannot exploit data declared much before 2017 because a major update of the tax administration's IT system in late 2016 disrupted the continuity of tax identification numbers, thus preventing us from identifying firms across the update. When studying the impact of the introduction of Etax, we also prefer not to exploit data beyond 2019, as the COVID-19 pandemic spread to Senegal in early 2020, affecting taxpayers differently depending on whether or not they had already adopted Etax before. Finally, we remove from the sample the 36 largest taxpayers who adopted e-taxation during the pilot phase starting in late 2016.

We define the treated group as all firms that have consistently been monitored by the LTC during the reference period. Figure 1 shows that first registrations and use took place in September 2017 and sharply increased afterwards. By January 2018, each month, 85 per cent of the taxpayers were using e-filing, and almost 80 per cent of them were using e-payment. The adoption took place quickly enough to allow for a credible impact evaluation, and by the end of the period, the e-filing rate was close to 100 per cent and the e-payment rate around 90 per cent. As a control group, we use all firms monitored by tax centres in the region of Dakar that did not adopt Etax - including MTC2 and RSC where a very small share of taxpayers did adopt Etax during the second half of 2018, probably on a voluntary basis (Figure 1). [5] The imperfect compliance with treatment is very minor and does not affect our results. For the sake of brevity, we do not show the results obtained under a fuzzy design. After this preliminary selection, the full sample consists of 5,559 taxpayers, 697 (12.5 per cent) of which are treated. By construction, treated and control taxpayers are very different. The main variable which the administration uses to allocate firms to the LTC is the turnover. To increase the comparability of the sample, we restrict our control group to its largest 703 taxpayers (measured by average pre-reform turnover). This generates our final dataset, with a total of 1,400 taxpayers.

[5] We exclude the MTC1 because adoption of Etax was progressive during 2018 and driven (at least partly) by self-selection.

4.3. Empirical strategy and estimation

Empirical strategy: We estimate the impact of e-taxation by comparing treated and control taxpayers before and after the reform. Treatment adoption is defined at the centre level, so the first treated period is the same across all taxpayers and, after this month, we consider *treated* any taxpayer monitored by the LTC, irrespective of whether it actually uses e-taxation technology. Our estimates thus capture the *intention-to-treat* impact of adopting e-taxation at the centre level on its taxpayers. Given these initial specifications, and following Roth et al., 2023 and Miller, 2023, we can estimate the impact of the reform using an event-study model with this functional form:

$$y_{it} = \alpha + \sum_{k=1}^{K} \beta^k \cdot \text{Lead}_{it}^k + \sum_{j=J}^{-2} \gamma^j \cdot \text{Lag}_{it}^j + \lambda_t + \delta_i + \epsilon_{it} \tag{1}$$

The statistical unit is the taxpayer (i) at the month level (t), with $J \leq t \leq -1$, or $1 \leq t \leq K$. Variables Lag_{it}^j and Lead_{it}^k are binary indicators equal to one for any observation from the treatment group taking place j months *before* the adoption of Etax and k months *after* respectively. Lead_{it}^1 indicates the first treated month, while Lag_{it}^{-1} (the last untreated month) is omitted from the equation 1 to serve as baseline. Each β^k captures the impact of the reform exactly *k-months* after its introduction compared to baseline ($t = -1$) and is equivalent to the β coefficient one can retrieve by estimating a simple difference-in-differences model with only two periods, either before treatment ($t = -1$) or after ($t = k$). The coefficients γ^j are similar but for the period *before* the reform, and act as placebo tests. The coefficients δ_i and λ_t are individual and time fixed-effects, and ϵ_{it} is the unobserved error term.

Next, to summarise into one single coefficient the overall impact for the entire post-reform period observed, we group all observations taking place after the reform and estimate the following 2×2 difference-in-differences model with fixed effects, where Treat_{it} and Post_{it} are indicators for being in the treated centre, and observed after the reform respectively.

$$y_{it} = \alpha + \beta \cdot \text{Treat}_{it} \cdot \text{Post}_{it} + \lambda_t + \delta_i + \epsilon_{it} \tag{2}$$

The underlying identification assumption requires that, absent the reform, outcomes y_{it} would have evolved similarly in the treatment and the control group. We systematically show placebo tests for our main outcome of interest on the pre-reform period and carefully discuss this issue in the results section. Provided that no other shock happened at the same time differently affecting treatment and control groups, this strategy allows us to attribute any notable changes observed around late 2017 as resulting from the reform. We discuss in appendix A.1 potential confounding factors related to staff reallocation.

Inference: Treatment takes place at the centre level. Accordingly, we should cluster our standard errors at the centre level (MacKinnon, Nielsen, and Webb, 2023). However, given that we have only ten centres (including the treated one), clustering at that level can produce implausibly low standard errors leading to severe over-rejections. In our case, even wild bootstrap can (but needs not) perform poorly (MacKinnon and Webb, 2018). Eventually, the best available method to quantify the uncertainty around our estimates is to perform randomisation inference for difference-in-differences as described and tested by MacKinnon and Webb, 2020. We use this method as our benchmark and systematically test wild bootstrap for robustness. To do so, we generate a distribution of our difference-in-differences estimators by repeatedly reassigning the treatment indicator according to the original design —preserving the number of treated centres (one) and the timing of treatment— and recomputing the estimates for each reassignment. The p-value for a two-sided null hypothesis is then computed as the proportion of placebo estimates whose absolute value is at least as large as that of the observed estimate, yielding exact, finite-sample inference that does not rely on asymptotic approximations.

5. Results

Our main results are displayed in Table 1 and Figures 2, 3 and 4. Following the recent econometric literature, we use randomisation inference as our main benchmark test for inference. We run both two-sided and one-sided significance tests and explicitly stress when one-sided is significant but against the hypothesis made in section 3. Yet, with only 10 clusters, the smallest possible p-value we may obtain with a randomisation inference test is 10 per cent, hence our design has arguably low power. In our discussion, we therefore also take stock of additional evidence on the dynamics of the treatment impact under wild bootstrap inference tests (Figure 2, 3 and 4) and raw trends (Figure A.2).

Our results suggest that the reform had a negative impact of about 2 percentage points on the probability of declaring during the first one or two months of the transition (September and October 2017). This is consistent with a fixed cost for both taxpayers and administrators to adapt to the new process. Estimates become positive soon after but not significantly different from zero (under neither randomisation inference nor wild bootstrap). However, the spread gradually increases over time, and by the very end of the period the probability of declaring is significantly larger (by 3 percentage points) for the treatment group under wild bootstrap and one-sided randomisation inference. Our results on the probability of declaring a positive liability follow a very similar trend (Figure A.3), but average liability declared did not increase nor decrease (Figure A.3). Such results are consistent with the expected impact described in section 3 according to which e-filing should allow tax officials to more easily detect missing declarations, or key information in declarations (e.g. positive liability) and thus rapidly enforce better reporting. However, given that we observe no rapid increase right after the introduction of e-filing, we cannot rule out that our estimates are biased by difference in cycles; the further we get from the month of the reform, the less plausible our parallel trends assumption becomes.

Next, we turn to the analysis of tax payments. The probability of paying some amount seemed to have decreased during the first year by about 4 percentage points. Monthly results are significant at the 5 per cent level under wild bootstrap method during that first year, but only inconsistently at the 10 per cent level for a one-sided test under randomisation inference. In 2019, the effect vanishes. This one-year negative effect is unexpected. Section 3 would, if anything, imply that the probability of paying would *increase*. One interpretation could be that some taxpayers momentarily resist the adoption of e-payment for fear of stronger enforcement later. Online payments indeed require taxpayers to share their bank identity with the administration, who can then use it to more easily activate a legal procedure forcing banks to remit any outstanding liability owed by a given taxpayer —provided there is sufficient provision in the account. Unfortunately, this is only speculative.

We find very limited results on average payment. The monthly estimates are positive (except for one month), and always significant at the 10 per cent level for one-sided tests under randomisation inference. Yet, they are significant at the 5 per cent level under wild bootstrap only in November and/or December (5 out of 26 post-treatment months). This pattern reflects a seasonal variation that should be interpreted with care. First, we do not observe these months in the pre-period —and therefore discard them in specification (4) of Table 1. Moreover, they could be the result of a confounding factor. Indeed, by the end of the year, the tax administration gets closer to its final performance results, which it has to report to the Ministry of Finance that evaluates it accordingly. November and December are typically months during which enforcement might increase if the administration is below its target, and tightening rules at the LTC specifically (rather than in all centres) may be an efficient strategic response, given the size of the taxpayers it monitors. Last and most importantly, we find aggregate evidence suggesting that the introduction of e-payment could have actually reduced the share of payments that were processed completely outside the information system. Although the practice is regarded as highly risky and allowed only in exceptional circumstances, it is still possible to process payment 'manually', that is without recording it into the information system. In that case, the taxpayers receive a paper proof of payment, and the administration records the transaction in an invoicing book. Paper-based payments are then added to digitised payments (including e-payment) when computing total tax collected. In Figure 3 (lowest panel, right), we plot the raw and smoothed difference between general accounting totals and aggregated micro data on payments, as a percentage of the former. We can make this comparison for PAYE data only. Despite monthly fluctuations, we observe a drop in the average residual share of payments that are not included in our micro data as online payment gets introduced. The volume of non-digitised payments is not trivial: in 2017, it was equal to USD 1,876,551 on average per month (excluding November and December). Given the estimate of column (4) in Table 1 and the number of treated firms, the volume of non-digitised PAYE payments alone could already account for a third of the average impact on payment, even though it does not account for non-digitised payment of VAT and WTS. Overall, our results show that the introduction of e-payment did not lead to a clear increase in average tax paid, but probably had a positive impact on the share of *digitised* payments. Combining information on liability declared and recorded payments, we also show that the introduction of e-taxation had no significant impact on average (positive) outstanding liability nor on average outstanding liability as a share of total liability (Figure A.3).

Finally, and in contrast with the above results, we find a sharp, large and persistent negative impact on the prevalence of missing values. The probability of a value being missing decreases by 7 percentage points (not shown here for brevity). Investigating exactly which values are not missing any more, we find that information from the PAYE declarations drives the results. In particular, the share of missing values in the field 'number of employees' drops by about 6 percentage points (Figure 4 and Table 1). The trends in Figure A.2 are consistent with what we would expect: missing values drop in the treated group and remain stable in the control group. Relative to pre-reform level, this represents a 90 per cent decrease. Beyond the short transition period following the reform, the month-specific coefficients are very stable and consistently significant at the 10 per cent level under one-sided test with randomisation inference, and at the 5 per cent level for two-sided tests under bootstrap method.

In light of this result, our interpretation is that the drop in missing values derives from the elimination of omissions made by data-entry agents. Values are mostly missing for fields that data-entry agents did not have to fill to validate a given declaration as long as they entered some value for the field 'tax liability'. As tax payments can still be processed as long as there is information about the tax liability, some level of tolerance for missing secondary information may have been granted to accelerate data-entry. Online declaration therefore increases data quality.[6] The implications of the increase in data quality are not trivial. Indeed, consistent with the sharp *drop* in missing value for the 'number of employees', the average number of employees observed sharply increases by 40 (see Table 1 and Figure 4). As for missing values, this result is very stable and significant throughout the period. This represents a 33 per cent rise relative to the baseline value in our treatment group. It is large and sudden enough for it to be salient even when plotting the monthly *aggregate* number of employees declared by all tax-registered firms in Senegal from 2014 to 2021 (Figure 4, lowest panel). In fact, the average increase we estimate from our sample of treated firms can account for two-thirds of the historic 30 per cent growth we observe in total formal employment between September and December 2017. In other words, due to missing values, total formal employment according to administrative data was underestimated by at least 20 per cent.

6. Conclusion

This paper studies the causal effects of mandatory electronic filing and payment on tax compliance among large taxpayers in Senegal. Exploiting the introduction of this technology in mid-2017 and using high-frequency administrative data in a dynamic difference-in-differences framework, we examine impacts on declaration behaviour, payments, and data quality.

[6] Incidentally, these results imply that any average impact whose measurement depends on information present in the declaration may become harder to interpret as there are now two sources of compositional effects: the composition of those who declare, and, conditional on declaring, that of those who do not have missing values. This is the reason why we do not include here impacts related to tax bases or tax rates.

The reform generated short-run transition costs: declaration rates temporarily declined immediately after implementation. Beyond this adjustment period, we find no robust evidence that e-filing or e-payment increased the probability of declaring, the liability declared, the probability of paying, or the average amount paid in the short to medium run. Any observed increase in tax payments among treated firms appears largely driven by improvements in the recording of transactions rather than changes in underlying compliance. However, the adoption of e-filing sharply and persistently reduced the prevalence of missing values in tax declarations. This effect reflects a change in the data production process—specifically, the elimination of omissions introduced during manual data-entry—rather than behavioural responses by taxpayers. These improvements have large measurement consequences. The transition to e-filing mechanically increased reported employment among treated firms, accounting for a sizeable share of the observed rise in aggregate formal employment after 2017. We show that missing information in pre-reform administrative data was leading to a substantial underestimation of formal employment.

This research contributes to the limited literature on the role of e-filing and e-payment technologies in low-income countries, offering nuanced insights into their effects on tax compliance. Overall, our results indicate that the introduction of digital tax filing and payment systems had limited effects on compliance outcomes for large taxpayers over the period studied, but fundamentally altered the quality of administrative data.

References

Ali, Merima, Abdulaziz Shifa, Abebe Shimeles, and Firew B Woldeyes (2015). “Information technology and fiscal capacity in a developing country: evidence from Ethiopia”. In.

Allingham, Michael G and Agnar Sandmo (1972). “Income tax evasion: A theoretical analysis”. In: *Journal of public economics* 1.3-4, pp. 323–338.

Almunia, Miguel and David Lopez-Rodriguez (2018). “Under the radar: The effects of monitoring firms on tax compliance”. In: *American Economic Journal: Economic Policy* 10.1, pp. 1–38.

Bachas, Pierre, Roberto N Fattal Jaef, and Anders Jensen (2019). “Size-dependent tax enforcement and compliance: Global evidence and aggregate implications”. In: *Journal of Development Economics* 140, pp. 203–222.

Besley, Timothy, Robin Burgess, Adnan Khan, and Guo Xu (2022). “Bureaucracy and development”. In: *Annual Review of Economics* 14, pp. 397–424.

Jouste, Maria, Milly I Nalukwago, and Ronald Waiswa (2021). *Do tax administrative interventions targeted at small businesses improve tax compliance and revenue collection? Evidence from Ugandan administrative tax data*. Tech. rep. WIDER Working Paper.

Luttmer, Erzo FP and Monica Singhal (2014). “Tax morale”. In: *Journal of economic perspectives* 28.4, pp. 149–168.

MacKinnon, James G, Morten Ørregaard Nielsen, and Matthew D Webb (2023). “Cluster-robust inference: A guide to empirical practice”. In: *Journal of Econometrics* 232.2, pp. 272–299.

MacKinnon, James G and Matthew D Webb (2018). “The wild bootstrap for few (treated) clusters”. In: *The Econometrics Journal* 21.2, pp. 114–135.

– (2020). “Randomization inference for difference-in-differences with few treated clusters”. In: *Journal of Econometrics* 218.2, pp. 435–450.

Milanez, Anna (2017). “Legal tax liability, legal remittance responsibility and tax incidence: Three dimensions of business taxation”. In: *OECD Taxation Working Papers* 32.

Miller, Douglas L (2023). “An Introductory Guide to Event Study Models”. In: *Journal of Economic Perspectives* 37.2, pp. 203–230.

OCDE (2022). “Revenue Statistics in Africa 1990-2020”. In.

Okunogbe, Oyebola and Victor Pouliquen (2022). “Technology, taxation, and corruption: evidence from the introduction of electronic tax filing”. In: *American Economic Journal: Economic Policy* 14.1, pp. 341–372.

Okunogbe, Oyebola and Fabrizio Santoro (2023). “The Promise and Limitations of Information Technology for Tax Mobilization”. In: *The World Bank Research Observer* 38.2, pp. 295–324.

Roth, Jonathan, Pedro HC Sant’Anna, Alyssa Bilinski, and John Poe (2023). “What’s trending in difference-in-differences? A synthesis of the recent econometrics literature”. In: *Journal of Econometrics*.

Santoro, Fabrizio, Razan Amine, and Tanele Magongo (2024). “Encouraging digital tax tools as a response to Covid: evidence from Eswatini”. In: *International Tax and Public Finance* 31.1, pp. 95–135.

Slemrod, Joel (2019). “Tax compliance and enforcement”. In: *Journal of Economic Literature* 57.4, pp. 904–954.

Slemrod, Joel and Tejaswi Velayudhan (2018). “Do firms remit at least 85% of tax everywhere? New evidence from India”. In: *Journal of Tax Administration* 4.1, pp. 24–37.

Weber, Max (1922). *Economy and Society: An Outline of Interpretive Sociology*. Vol. 1. University of California Press.

World Bank, World (2016). *World development report 2016: Digital dividends*. World Bank Publications.

Figure 1 – Adoption of Etax - Monthly share of effective users

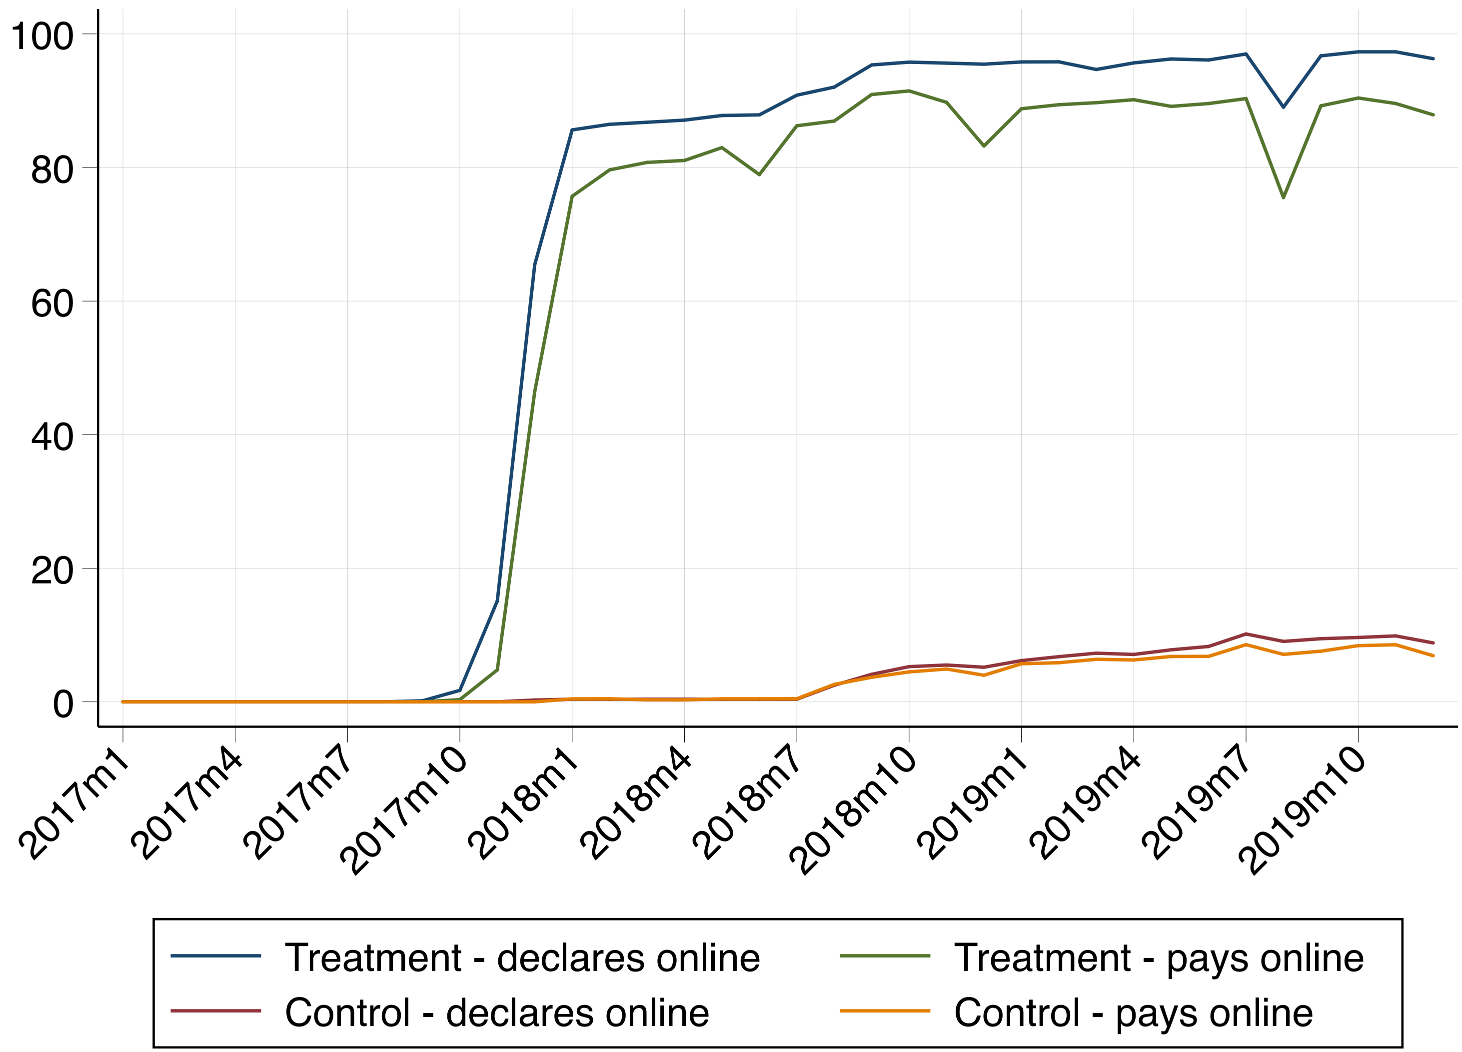

Notes. The figure displays the share of taxpayers that effectively use Etax each month, either through e-filing or e-payment, within each group. The sample here is the sample we use for all our main estimations, with 697 treated taxpayers and 703 taxpayers in the control group.

Figure 2 – Impact on declaration

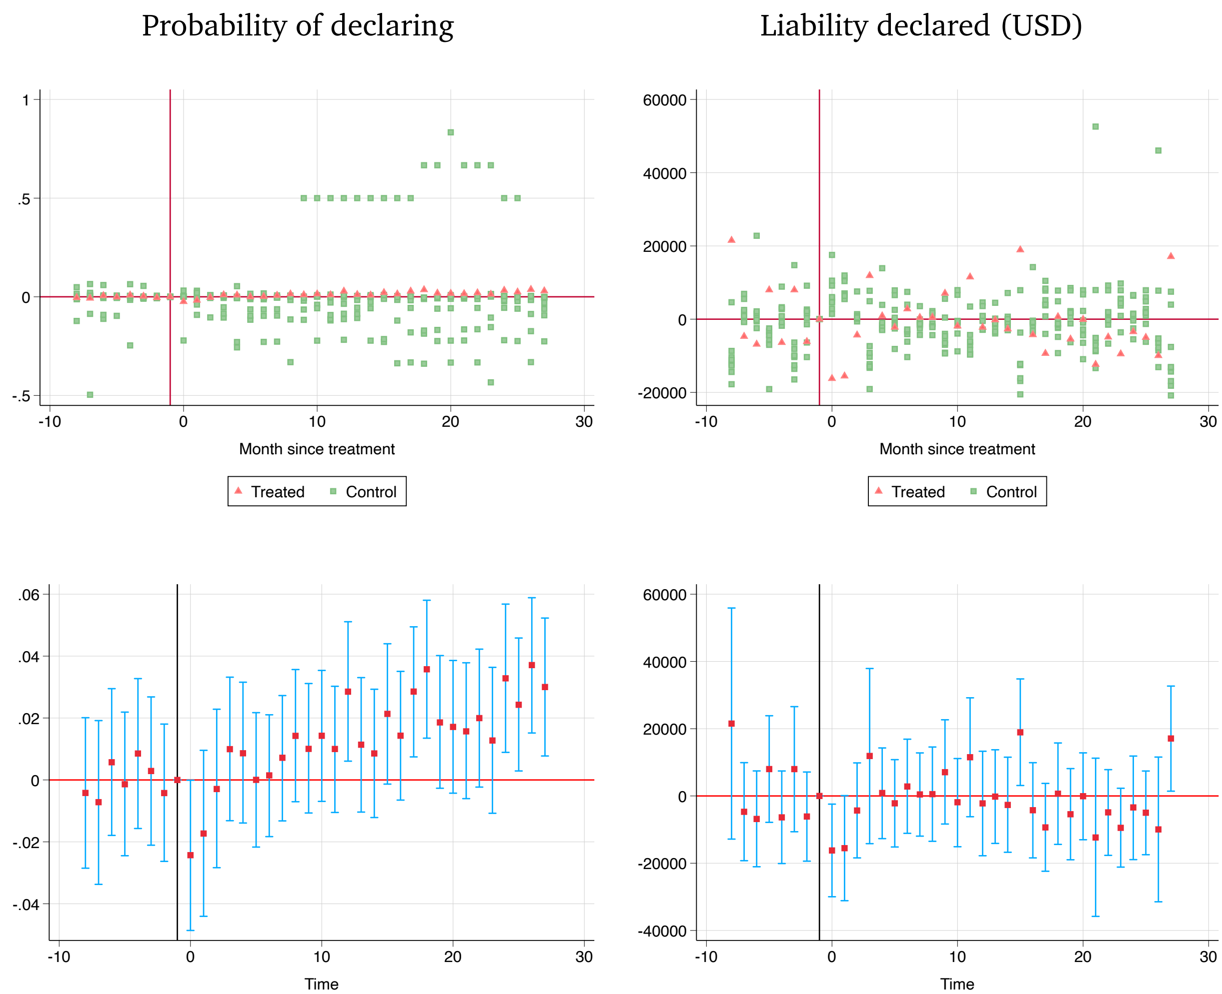

Notes: The upper panel plots the month specific distribution of β^k coefficients we estimate with equation 1, by assigning the treatment status to each of the different tax unit (i.e. clusters) in our sample. We highlight in red the coefficients retrieved when the tax unit assigned to treatment is the actually treated one (i.e. Large Taxpayers Unit), while green estimates are placebos specifications. The lower panel plots event study estimates of 1 with 95 per cent interval derived with wild bootstrap method.

Figure 3 – Impact on tax payments

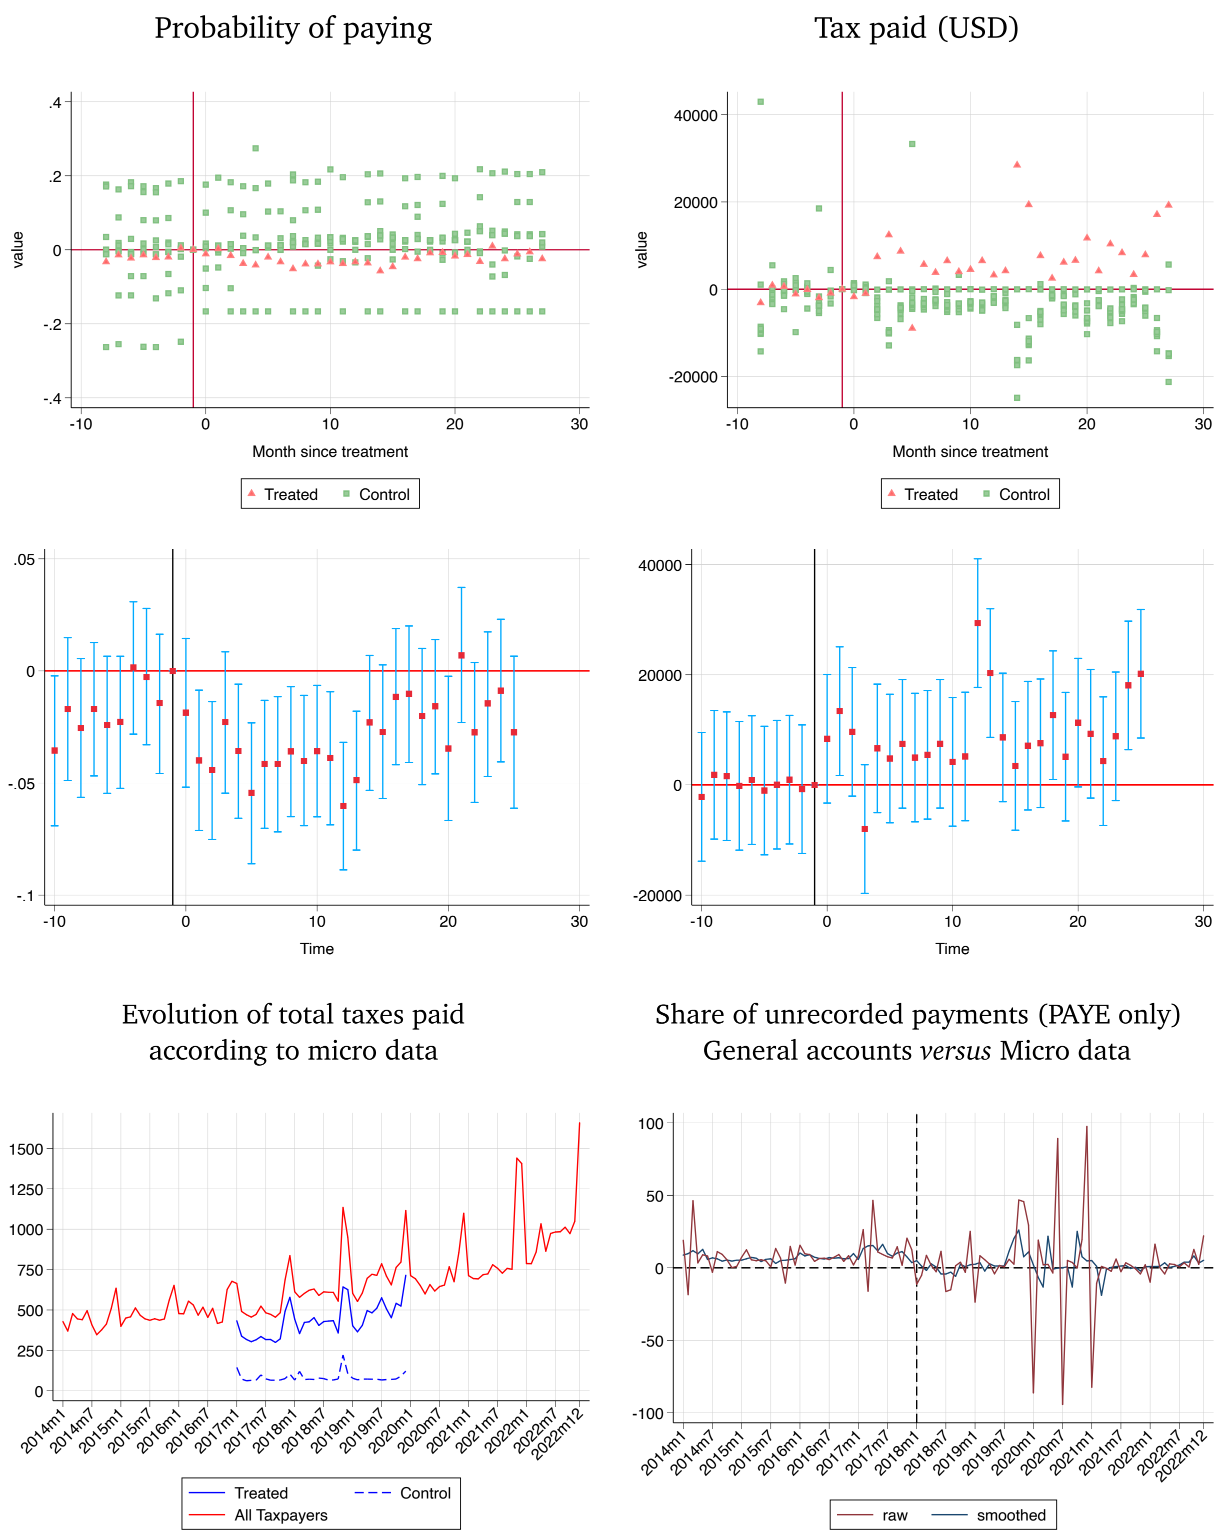

Notes: The *upper panel* plots the month specific distribution of β^k coefficients we estimate with equation 1, by assigning the treatment status to each of the different tax unit (i.e. clusters) in our sample. We highlight in red the coefficients retrieved when the tax unit assigned to treatment is the actually treated one (i.e. Large Taxpayers Unit), while green estimates are placebos specifications. The *middle panel* plots event study estimates of 1 with 95 per cent interval derived with wild bootstrap method clustering at the tax unit level. In the *lowest panel*, the left figure plots the monthly evolution of aggregate payments based on the universe of digitised payment data. The *last figure* displays the raw and smoothed evolution of the difference between total PAYE payments recorded in the aggregate general accounts and the total PAYE payment recorded in the digitised micro data, as a share of the first term. The PAYE tax is the only tax for which this comparison can be made.

Figure 4 – Impact on missing values and their measurement implications

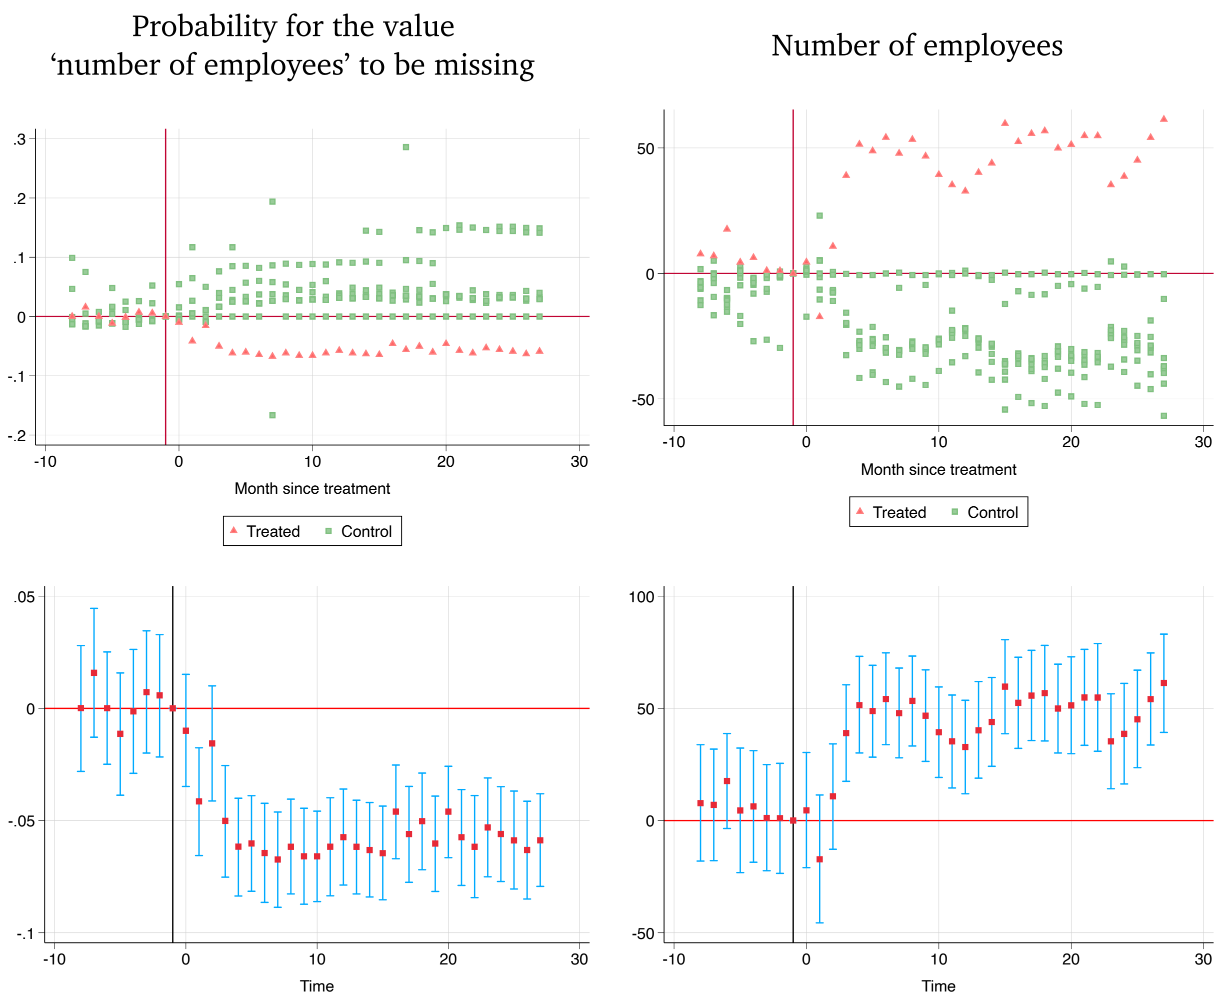

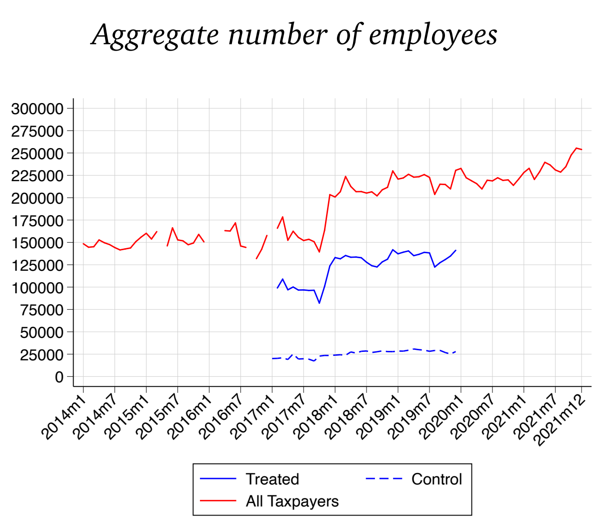

Notes: The *upper panel* plots the month specific distribution of β^k coefficients we estimate with equation 1, by assigning the treatment status to each of the different tax unit (i.e. clusters) in our sample. We highlight in red the coefficients retrieved when the tax unit assigned to treatment is the actually treated one (i.e. Large Taxpayers Unit), while green estimates are placebos specifications. The *middle panel* plots event study estimates of equation 1 with 95 per cent interval derived from wild bootstrap method clustering at the tax unit level. The *last figure* plots aggregate trends for the number of employees, estimated by summing information declared by employers who remit the PAYE —we leave as missing the months for which too few declarations were digitised.

Table 1 – Synthetic Results

Statistics		Declares	Liability declared	Pays	Tax paid	The number of employees is missing	Number of employees
		(binary)	(cont.)	(binary)	(cont.)	(binary)	(cont.)
		(1)	(2)	(3)	(4)	(5)	(6)
β		0.02	-1907.45	-0.01	6186.11	-0.06	41.01
two-sided p-value	Random Inf.	0.60	0.50	0.80	0.20	0.30	0.10
	Wild Boot.	0.10	0.17	0.32	0.13	0.12	0.07
one-sided p-value	Random Inf.	0.30	0.70	0.10	0.10	0.10	0.10
	Wild Boot.	0.05	0.91	0.16	0.07	0.07	0.04
Firm f-e		Y	Y	Y	Y	Y	Y
Time f-e		Y	Y	Y	Y	Y	Y
# tax units		10	10	10	10	10	10
# firms		1,400	1,400	1,400	1,400	1,400	1,400
N		46,200	46,200	42,000	42,000	46,200	46,200

Notes. This table presents estimation results of equation 2 for all our main outcomes. The first line displays the coefficients we obtain for the (Treat X Post) variable and measures the impact for the full post-treatment period. The subsequent lines compare the p-values we obtain under different test options: randomisation inference (our benchmark as advocated in MacKinnon and Webb, 2020) and wild bootstrap (9,999 replications and a correction for small number of clusters). By design, with only 10 clusters, the smallest p-value we can obtain with randomisation inference is 10 per cent. Specifications (2) and (4) set the first month of treatment as November 2017 when firms start adopting online payment, while other specification set the first month of treatment as September 2017, the start of online filing. We remove the months of November and December under specification (3) and (4) as discussed in section 5.

A. Appendix

A.1. On the potential reallocation of administrative staff

In anticipation of the reform, or as an immediate response to it, it is possible that the tax administration decided to relocate staff across treatment and control groups. This may have been a natural strategy considering that online services are adopted at least partly to save the administration the cost of processing papers. In that regard, the staff occupied with digitising and classifying declarations ex-ante at the LTC may have been relocated to other (control) centres. Movements in the other direction would also have been possible, whereby high-skilled staff were relocated from control centres to the LTC. In either case, this could represent a threat to our identification strategy.

Compiling data on the yearly evolution of the distribution of all civil servants working at the tax administration, we show that none of these changes have occurred. Table A.1 shows that changes between 2017 and 2018 in the number of civil servants allocated to the treatment centre were very marginal: +2 among civil servants of rank A (highest qualification), -1 for rank B, +1 for rank C and no change for the lowest rank D. Simultaneously, changes observed for civil servants allocated to centres of the control group mostly affected the most qualified sub-group: +13 and +19 for ranks A and B respectively (i.e. 20 and 12 per cent increase respectively), while rank C and D civil servants remain virtually constant. Centres from the control group received 13 (i.e. 70 per cent) of the 19 newly recruited rank A civil servants. Overall, we observe that changes in the staffing structure have been minor and are unlikely to drive our results. If anything, we would only expect the increase in rank A civil servants among centres in the control group to make our estimates slightly downward biased.

Table A.1 – Number of Civil Servants by treatment status and qualification - 2016-2022

		2016	2017		2018		2019		2020		2021		2022	
		N.	N.	% change	N.	% change	N.	% change	N.	% change	N.	% change	N.	% change
A) Treatment center (LTU)	A	34	34	0	36	5.8	34	-5.5	35	2.9	29	-17.1	30	3.4
	B	28	31	10.7	30	-3.2	29	-3.3	29	0	37	27.5	38	2.7
	C	13	13	0	14	7.6	14	0	14	0	13	-7.1	13	0
	D	4	4	0	4	0	3	-25	3	0	2	-33.3	0	-100
	Total	79	82	3.7	84	2.4	80	-4.7	81	1.2	81	0	81	0
B) Control centers	A	67	66	-1.4	79	19.7	94	18.9	92	-2.1	97	5.4	96	-1.0
	B	167	165	-1.2	184	11.5	193	4.8	193	0	216	11.9	216	0
	C	241	253	4.9	252	-0.4	249	-1.1	249	0	266	6.8	257	-3.3
	D	33	37	12.1	35	-5.4	33	-5.7	33	0	35	6.0	35	0
	Total	508	521	2.5	550	5.5	569	3.4	567	-0.3	614	8.2	604	-1.6
C) Other centers/departments	A	159	148	-6.9	152	2.7	168	10.5	169	0.6	187	10.6	195	4.2
	B	218	220	0.9	211	-4.0	220	4.2	220	0	218	-0.9	218	0
	C	166	169	1.8	169	0	170	0.5	170	0	187	10	190	1.6
	D	71	69	-2.8	65	-5.8	70	7.6	70	0	61	-12.8	61	0
	Total	614	606	-1.2	597	-1.4	628	5.1	629	0.1	653	3.8	664	1.6
D) All	A	260	248	-4.6	267	7.6	296	10.8	296	0	313	5.7	321	2.5
	B	413	416	0.7	425	2.1	442	4	442	0	471	6.5	472	0.2
	C	420	435	3.5	435	0	433	-0.4	433	0	466	7.6	460	-1.2
	D	108	110	1.8	104	-5.4	106	1.9	106	0	98	-7.5	96	-2.0
	Total	1201	1209	0.6	1231	1.8	1277	3.7	1277	0	1348	5.5	1349	0.0

Notes. This table displays the number of civil servants - by qualification - allocated to each quasi-experimental group and its evolution since 2016. Panel (A) consists of civil servants allocated to the Large Taxpayers Centre (the treated centre in our main specification); panel (B) consists of civil servants affected to centres which are included in our control group (all centres of Dakar, excluding MTC1); panel (C) gathers all other civil servants - this includes regional centres as well as all staff working at function-specific centralised departments such as the IT, legislation, human resources. Finally panel (D) is the sum of (A),(B) and (C).

A.2. On synthetic control

In the specific case of this paper, using the synthetic control method at the tax centre level is not a more appropriate option for inference than the event study we implement in combination with randomisation inference or bootstrap. Indeed, for most outcomes, the average levels before the reform in the treated centre are higher than in any of the control centres, therefore the treated unit cannot be recovered using a convex combination of the control units. In that regard, the difference-in-differences method better allows us to cancel out level discrepancies in pre-trend levels. We illustrate this for our main result in Figure A.1.

Figure A.1 – Comparing Event-Study with Synthetic Control Method
“Probability for the ‘number of employees’ to be missing”

A) Raw trends

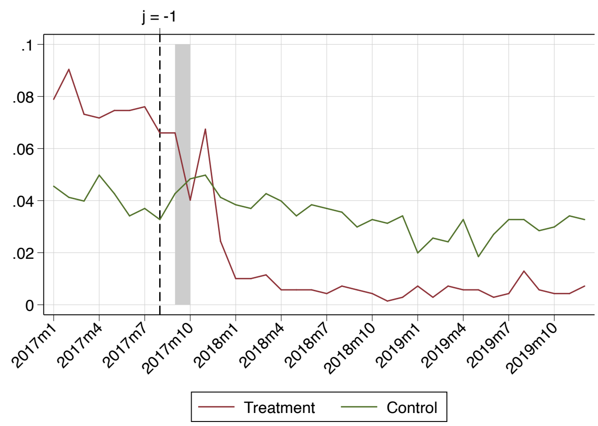

B) Synthetic control

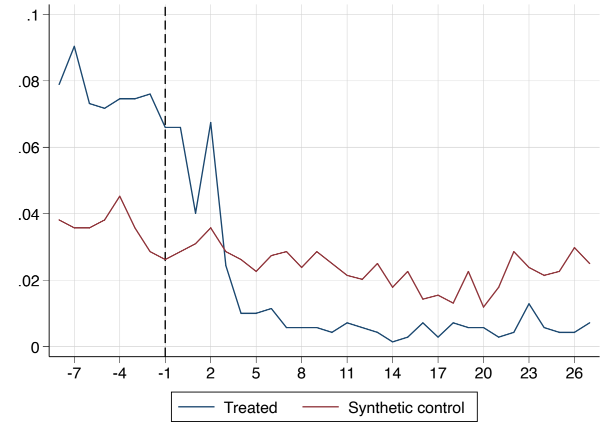

C) ... removing firm fixed-effects

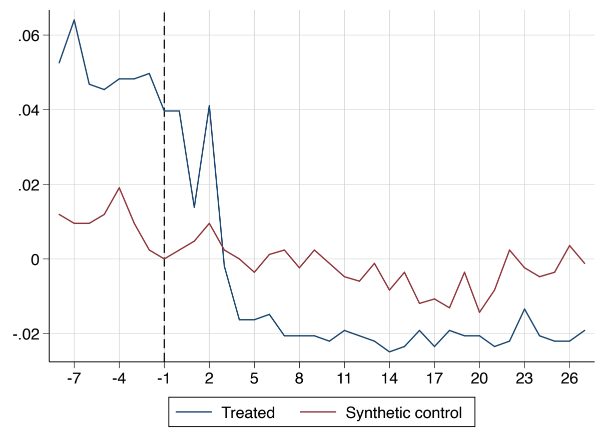

D) ... relative to $t = -1$

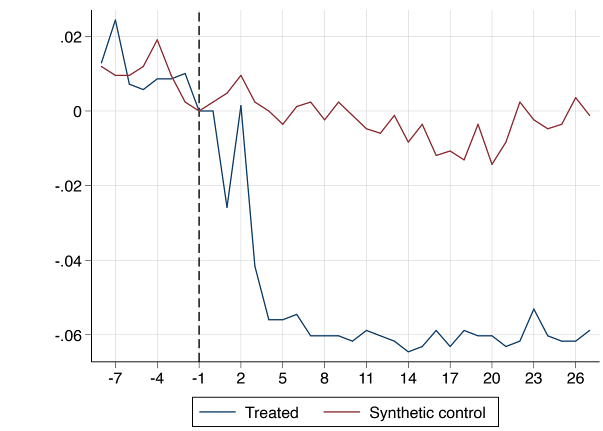

E) ... implied gap

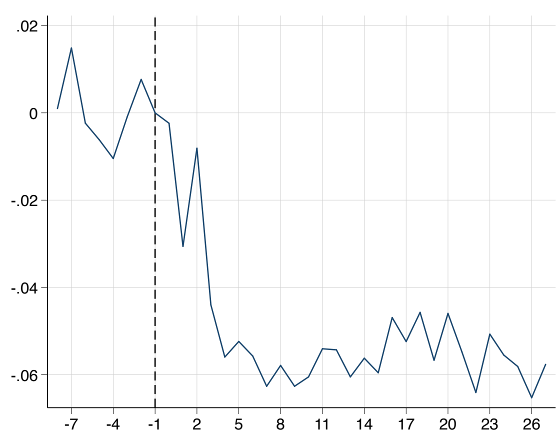

F) Event study with wild bootstrap inference

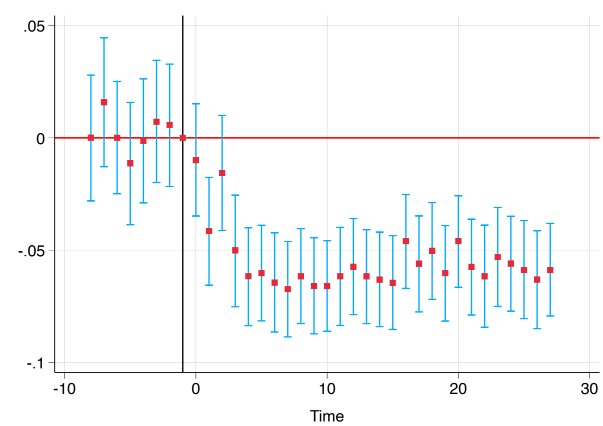

Notes: The first and last sub-figures reproduce Figures A.2.E and 4 (middle panel, left). This allows the reader to graphically compare those results with the ones obtained with synthetic control method. Panel B) compares treated and (simple) synthetic method (using only pre-treatment outcome values as predictors). Panel C) does the same but on the residuals obtained after having removed firms’ fixed-effect. Panel D) takes the difference relative to the value in $t = -1$ and Panel E) displays the gap between “Treated” and “Synthetic Control” of panel D). We observe that simple synthetic control method does a poor job at matching pre-trends. However, our event study results are very consistent in magnitude and trends with a synthetic difference-in-differences method on residuals.

A.3. Additional Figures and Tables

Figure A.2 – Raw trends

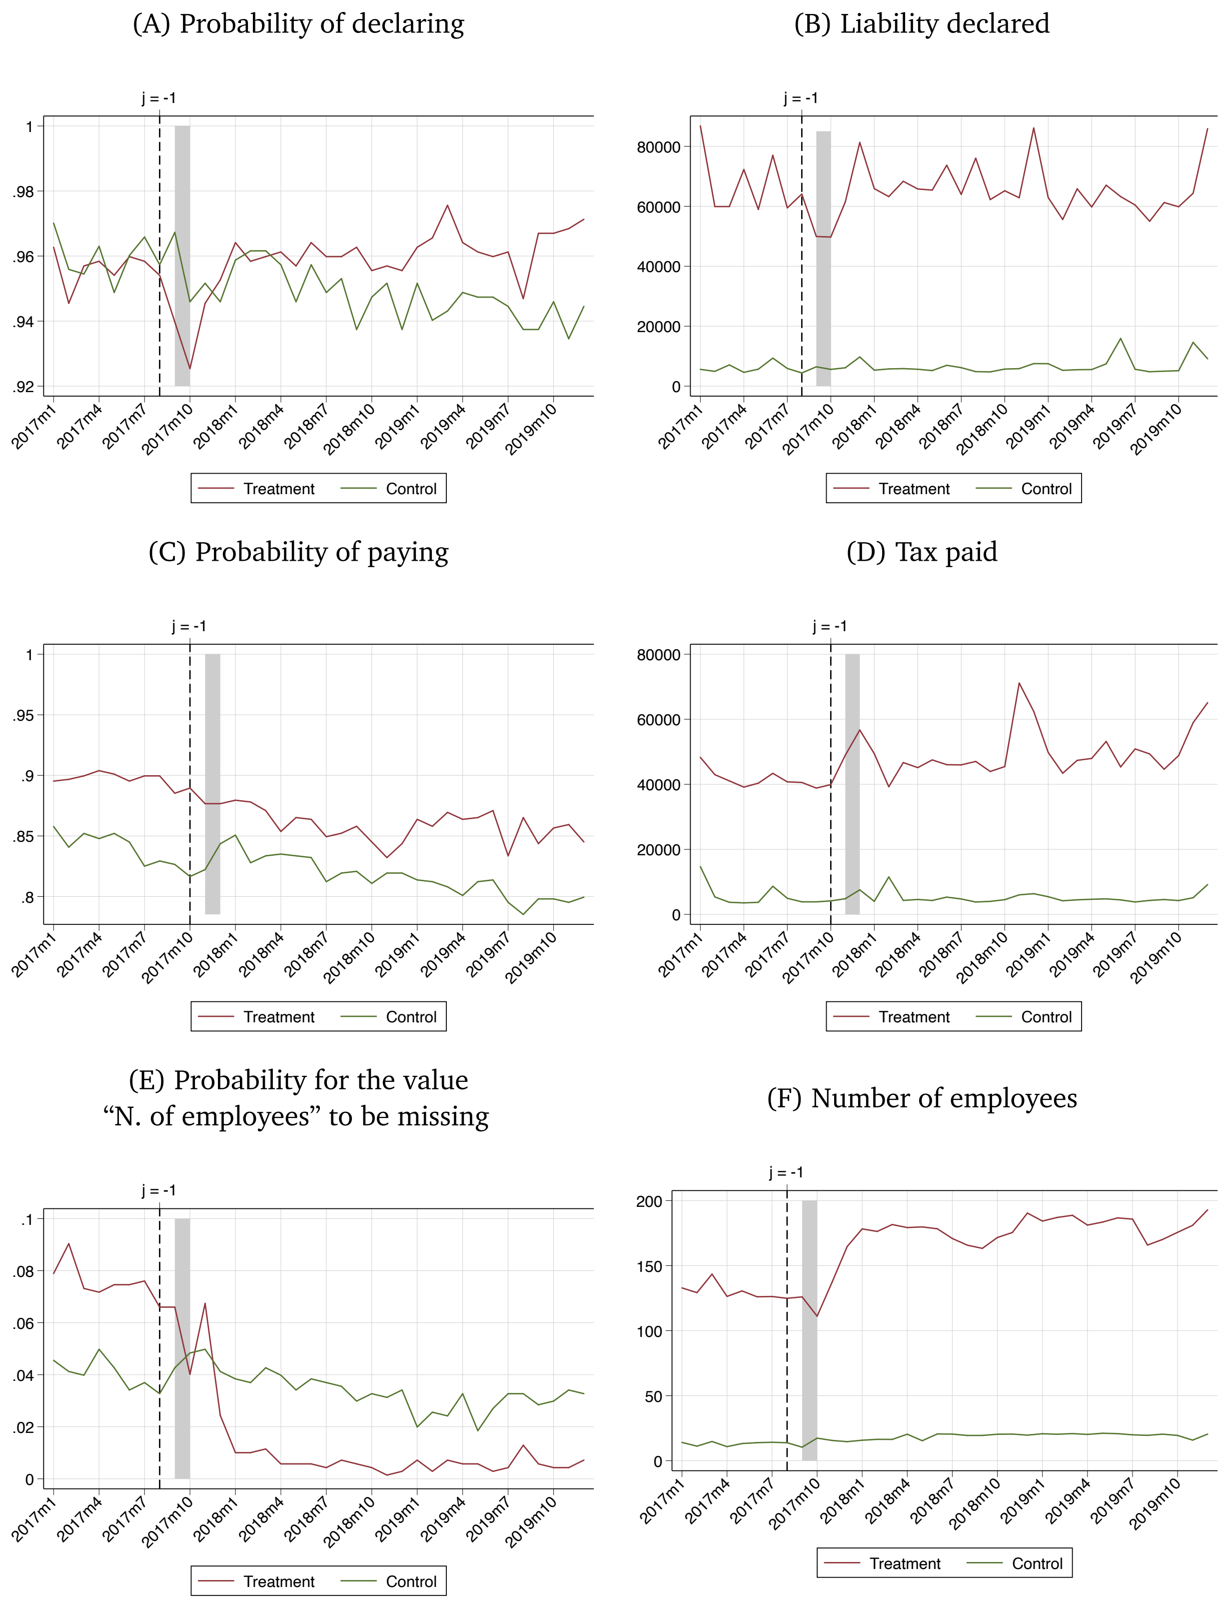

Notes: This figure displays raw trends for our six main outcomes. Treatment start ($t = 0$) is November 2017 for outcomes measured with information derived from payment data (C and D), and September 2017 for outcomes measured with information derived from tax declarations (A, B, E and F).

Figure A.3 – Additional results on (positive) outstanding liability

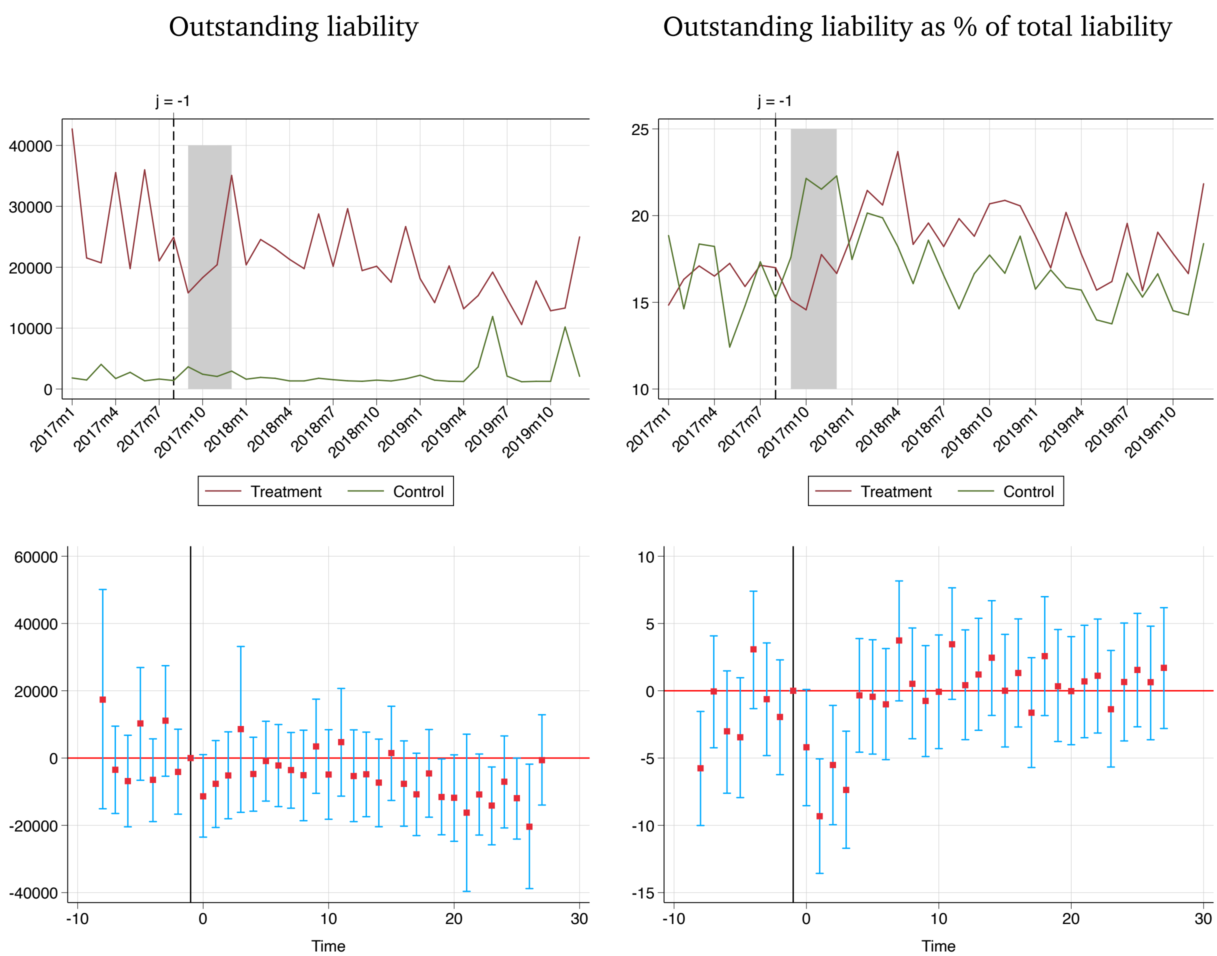

Notes: This figure displays raw trends and event study estimates (β^k coefficients of equation 1) for the additional outcome variables combining both declarative and payment data. Outstanding liability is the *positive* difference between total tax liability (VAT, PAYE and WTS — see section 4). We restrict to positive values only as, because of the possibility for some declarations to be missing, negative values can be misleading. Treatment start ($t = 0$) is September 2017. The grey area highlights the transition period —4 months as we use both declarative and payment data. Standard errors are clustered at the tax centre level using wild bootstrap method.